

CLAIMS

1. A method for generating an installation file for a particular version of a relational database comprising:

automatically determining a first set of data definition language (DDL) scripts associated with the particular version of the relational database;

automatically determining a second set of data manipulation language scripts associated with the particular version of the relational database; and

generating an installation file comprising a union of the first set and the second set.

2. The method as recited in claim 1 wherein the particular version is associated with a first version in a sequence of one or more versions of the relational database.

3. The method as recited in claim 2 wherein the automatically determining a first set comprises extracting a filename from metadata associated with the first version, the filename associated with a file comprising a data definition language script.

4. The method as recited in claim 2 wherein the automatically determining a second set comprises extracting a filename from metadata associated with the first version, the filename associated with a file comprising a data manipulation language script.

5. The method as recited in claim 2 wherein the generating an installation file comprises copying a data definition language script from a script file associated with the first set into the installation file.

6. The method as recited in claim 2 wherein the generating an installation file comprises copying a data manipulation language script from a script file associated with the second set into the installation file.

7. The method as recited in claim 6 wherein the copying further comprises prepending a create command to the data manipulate language script in the installation file.

8. The method as recited in claim 1 wherein metadata exists that describes a sequence of multiple versions of the relational database where each version is an upgrade from a previous version, and the particular version is not a first version in the sequence.

9. The method as recited in claim 8 wherein the metadata comprises an XML file.

10. The method as recited in claim 8 wherein the automatically determining a first set comprises:

extracting a set A_1 comprising one or more filenames from metadata associated with a first version in the sequence, the one or more filenames associated with a file comprising a data definition language script associated with the first version;

iteratively extracting a set A_i comprising zero or more filenames from metadata associated with an i^{th} version of the relational database, the zero or more filenames each associated with a file comprising a data definition language script to be executed when upgrading from version $i-1$ of the relational database to version i of the relational database, where i varies incrementally from 2 to j , where the particular version is j ;

determining the first set as the union of sets A_1, A_2, \dots, A_j .

11. The method as recited in claim 8 wherein the automatically determining a second set comprises:

extracting a set A_1 comprising one or more filenames from metadata associated with a first version in the sequence, the one or more filenames associated with a file comprising a data manipulation language (DML) script associated with the first version;

iteratively extracting a set A_i comprising zero or more filenames from metadata associated with an i^{th} version of the relational database, the zero or more filenames each associated with a file comprising a DML script to be executed to add or modify a DML object when upgrading from version $i-1$ of the relational database to version i of the relational database, where i varies incrementally from 2 to j , where the particular version is j ;

iteratively extracting a set B_i comprising zero or more filenames from metadata associated with an i^{th} version of the relational database, the zero or more filenames each associated with a file comprising a DML drop script to be executed to drop a DML object when upgrading from version $i-1$ of the relational database to version i of the relational database, where i varies incrementally from 2 to j , where the particular version is j ; and

determining the second set C_j by determining:

$$C_2 = [A_1 \cup A_2] - B_2,$$

$$C_3 = [C_2 \cup A_3] - B_3,$$

$$C_4 = [C_3 \cup A_4] - B_4,$$

...

$$C_j = [C_{j-1} \cup A_j] - B_j.$$

12. One or more computer-readable media having computer-readable instructions thereon which, when executed by a computer, cause the computer to implement the method as recited in claim 1.

13. A method for generating an upgrade file to upgrade version i of a relational database to version j of the relational database, where $j > i$, the method comprising:

determining a set A of data definition language (DDL) scripts that, when executed, perform creates, alters, and drops of DDL objects associated with version i of the relational database, resulting in DDL objects associated with version j of the relational database;

determining a set B of data manipulation language (DML) scripts that, when executed, create DML objects that are associated with version j of the relational database, but that are not associated with version i of the relational database;

determining a set C of DML scripts that, when executed, modify DML objects that are associated with both version i and version j of the relational database, but that differ between version i and version j of the relational database;

determining a set D of DML drop scripts that, when executed, drop DML objects that are associated with version i of the relational database, but that are not associated with version j of the relational database; and

generating an upgrade file comprising a union of sets A, B, C, and D ($A \cup B \cup C \cup D$).

14. The method as recited in claim 13 wherein the determining a set A comprises:

iteratively extracting sets M_k , each comprising zero or more filenames from metadata associated with a k^{th} version of the relational database, where $i < k \leq j$, the zero or more filenames each associated with a file comprising a data definition language script to be executed when upgrading from version $k-1$ of the relational database to version k of the relational database; and

determining the set A as the union of sets M_{i+1} , M_{i+2} , ..., M_j
($A = M_{i+1} \cup M_{i+2} \cup \dots \cup M_j$).

15. The method as recited in claim 13 wherein the determining a set B comprises:

determining a set E of DML scripts that when executed:

perform alters of DML objects associated with version i and version j of the relational database, but that differ between version i and version j of the relational database; and

perform creates of DML objects that are associated with version j of the relational database but that are not associated with version i of the relational database; and

determining set B as the difference between sets E and C ($B = E - C$).

16. The method as recited in claim 15 wherein the determining a set E comprises:

iteratively determining a set P_x of DML scripts that when executed will upgrade DML objects from version $x-1$ of the relational database to version x of the relational database, where x varies incrementally from $i+1$ to j ;

iteratively determining a set N_x of DML scripts that when executed will drop DML objects that are associated with version $x-1$ of the relational database but that are not associated with version x of the relational database, where x varies incrementally from $i+2$ to j ;

iteratively determining a set M_x of DML scripts that when executed will upgrade DML objects from version i of the relational database to version x of the relational database, where x varies incrementally from $i+1$ to j , and where:

$$M_{i+1} = P_{i+1}$$

$$M_{i+2} = [M_{i+1} \cup P_{i+2}] - N_{i+2}$$

$$M_{i+3} = [M_{i+2} \cup P_{i+3}] - N_{i+3}$$

...

$$M_j = [M_{j-1} \cup P_j] - N_j; \text{ and}$$

determining set $E = M_j$.

17. The method as recited in claim 13 wherein the determining a set C comprises:

determining a set E of DML scripts that when executed:

perform alters of DML objects associated with version *i* and version *j* of the relational database, but that differ between version *i* and version *j* of the relational database; and

perform creates of DML objects that are associated with version *j* of the relational database but that are not associated with version *i* of the relational database;

determining a set F_j of DML scripts that when executed, create DML objects associated with version *j* of the relational database; and

determining set C as the intersection of set E and set F_j ($C = E \cap F_j$).

18. The method as recited in claim 17 wherein the determining a set F_j comprises:

extracting a set M_1 comprising one or more filenames from metadata associated with a first version in the sequence, the one or more filenames associated with a file comprising a data manipulation language (DML) script associated with the first version;

iteratively extracting a set M_x comprising zero or more filenames from metadata associated with version x of the relational database, the zero or more filenames each associated with a file comprising a DML script to be executed to add or modify a DML object when upgrading from version $x-1$ of the relational database to version x of the relational database, where x varies incrementally from 2 to j ;

iteratively extracting a set B_x comprising zero or more filenames from metadata associated with version x of the relational database, the zero or more filenames each associated with a file comprising a DML drop script to be executed to drop a DML object when upgrading from version $x-1$ of the relational database to version x of the relational database, where x varies incrementally from 2 to j ;
and

determining the set F_j by determining:

$$F_2 = [M_1 \cup M_2] - B_2,$$

$$F_3 = [F_2 \cup M_3] - B_3,$$

$$F_4 = [F_3 \cup M_4] - B_4,$$

...

$$F_j = [F_{j-1} \cup M_j] - B_j.$$

19. The method as recited in claim 13 wherein the determining a set D comprises:

determining a set E of DML scripts that when executed:

perform alters of DML objects associated with version i and version j of the relational database, but that differ between version i and version j of the relational database; and

perform creates of DML objects that are associated with version j of the relational database but that are not associated with version i of the relational database;

iteratively determining a set F_x of DML scripts that when executed, drop DML objects associated with version $x-1$ of the relational database that are not associated with version x of the relational database, where x varies incrementally from $i+1$ to j ;

determining a set G as the union of sets $F_i, F_{i+1}, F_{i+2}, \dots, F_j$ ($G = F_i \cup F_{i+1} \cup \dots \cup F_{i+2}$); and

determining set D as the difference between set G and set E ($D = G - E$).

20. A system comprising:

one or more data definition language (DDL) scripts, each associated with one or more versions of a relational database;

one or more data manipulation language (DML) scripts, each associated with one or more versions of the relational database;

a database schema version management structure definition;

schema data associated with multiple versions of the relational database, the schema data organized according to the database schema version management structure definition; and

an installation file generator configured to apply laws of set theory to the schema data to generate a file comprising the one or more DDL scripts associated with a particular one of the multiple versions of the relational database, and the one or more DML scripts associated with the particular one of the multiple versions of the relational database.

21. The system as recited in claim 20 wherein the database schema version management structure definition comprises an XML schema definition.

22. The system as recited in claim 21 wherein the schema data is maintained in an XML file structured according to the XML schema definition.

23. A system comprising:

A memory;

A processor; and

a database schema version management system stored in the memory, executed on the processor, and configured to manage schema data associated with multiple versions of a relational database.

24. The system as recited in claim 23 wherein the schema data identifies a script associated with a data definition language object of the relational database.

25. The system as recited in claim 23 wherein the schema data identifies a script associated with a data manipulation language object of the relational database.

26. The system as recited in claim 23 wherein the database schema version management system is further configured to generate an installation file associated with an initial version of the relational database.

27. The system as recited in claim 23 wherein the database schema version management system is further configured to generate an installation file associated with a non-initial version of the relational database.

28. The system as recited in claim 23 wherein the database schema version management system is further configured to generate an upgrade file for upgrading a first version of the relational database to another version of the relational database.

29. One or more computer-readable media comprising computer-readable instructions which, when executed, cause a computer system to:

maintain schema data that identifies scripts associated with database objects of multiple sequential versions of a relational database; and

generate an installation file associated with an initial version of the relational database by applying laws of set theory to the schema data to identify scripts associated with the database objects of the initial version of the relational database.

30. The one or more computer-readable media as recited in claim 29, further comprising computer-readable instructions which, when executed, cause a computer system to:

generate an installation file associated with a non-initial version of the relational database by applying laws of set theory to the schema data to identify:

scripts associated with data definition language (DDL) objects that are associated with the non-initial version of the relational database; and

scripts associated with data manipulation language (DML) objects that are associated with the non-initial version of the relational database.

31. The one or more computer-readable media as recited in claim 29, further comprising computer-readable instructions which, when executed, cause a computer system to:

generate an upgrade file associated with an upgrade from a first, but not necessarily initial, version of the relational database to a second, later, but not necessarily immediately sequential, version of the relational database by applying laws of set theory to the schema data to identify:

data definition language (DDL) scripts associated with DDL objects of the database that have been created or modified between the first and second versions of the relational database;

data manipulation language (DML) scripts associated with DML objects of the database that have been created between the first and second versions of the relational database;

DML scripts associated with DML objects of the database that have been modified between the first and second versions of the relational database; and

drop scripts associated with database objects that have been dropped and not re-created between the first and second versions of the relational database.